arXiv:1610.07874v1 [math.PR] 25 Oct 2016

ROBUSTNESS OF MIXING UNDER ROUGH ISOMETRY, VIA BOTTLENECK SEQUENCES

LOUIGI ADDARIO-BERRY AND MATTHEW I. ROBERTS

Abstract. We provide new upper bounds for mixing times of general finite Markov chains. We use these bounds to show that the total variation mixing time is robust under rough isometry for bounded degree graphs that are roughly isometric to trees.

1. Introduction

The mixing time is one of the most fundamental and well-studied quantities associated to a Markov chain. It is natural to ask how robust the mixing time is: how much can we can change the mixing time by making small changes to the chain?

It has recently been shown, by Ding and Peres [5] and by Hermon [9] respectively, that neither total variation mixing times nor uniform mixing times are geometrically robust: in general, bounded perturbations of edge conductances can change both by arbitrarily large factors. For the moment, there is no general recipe which determines whether or not a given collection of graphs is robust in this sense.

One of the aims of this paper is to show that for a large class of chains—roughly speaking, any chain whose underlying graph is globally tree-like—the total variation mixing time is geometrically robust.

In order to do this we introduce new upper bounds on the mixing time that may be useful in their own right. It is well-known that the mixing time is related to bottlenecks in the graph, and our main idea is that the key quantity is the number and strength of bottlenecks that can be lined up in a row. We define a *bottleneck sequence* to quantify this concept and use it heavily throughout the article.

The main results of this paper are Theorem 1.1, which is a statement about robustness of mixing times; Theorem 1.2, which gives an upper bound on mixing times using bottleneck sequences; and Theorem 1.3, which provides a stronger but more complicated upper bound on mixing times involving a game between two players.

1.1. Notation.

Let $X = (X_n, n \geq 0)$ be an irreducible Markov chain on a finite state space V, and for $u, v \in V$ let $p_{uv} = \mathbb{P}_u(X_1 = v)$, where $\mathbb{P}_u$ signifies that the chain starts from u. Let E be the set of pairs $\{u, v\}$ with $u \neq v$ such either $p_{uv} > 0$ or $p_{vu} > 0$, and let $G = (V, E)$. Occasionally we will write G_X to signify that the graph G corresponds to the Markov chain X. To avoid periodicity issues, we assume throughout that X is *lazy*, i.e. that $p_{vv} = 1/2$ for all $v \in V$. Write $\pi = (\pi(v), v \in V)$ for the stationary distribution of X.

We say that X is *ε-uniform* if $\varepsilon\pi(x)p_{xy} \leq \pi(u)p_{uv} \leq \pi(x)p_{xy}/\varepsilon$ for all $uv, xy \in E$.

If X is reversible, i.e. $\pi(x)p_{xy} = \pi(y)p_{yx}$ for all $xy \in E$, then we write $c_{xy} = \pi(x)p_{xy}$ for the *conductance* of the edge xy.

A fundamental property of X is the *total variation mixing time*,

$$t_{\mathrm{mix}}(X) = \min\{n \geq 0 : \sup_v \|\mathbb{P}_v(X_n \in \cdot) - \pi\| \leq 1/4\},$$

where $\|\mu - \nu\| = \max_{A \subset V} |\mu(A) - \nu(A)|$ denotes the total variation distance between probability measures μ and ν on V.

A *stopping rule* for our Markov chain is a stopping time that is allowed to use extra randomness, as well as the history of the chain, to decide when to stop. For $v \in V$, we say that τ is a stopping rule from v to π if τ is a stopping rule and $\mathbb{P}_v(X_\tau = u) = \pi(u)$ for each $u \in V$. We define two more quantities that measure how long the chain takes to mix,

$$t_{\mathrm{stop}}(X) = \max_{v \in V} \inf\{\mathbb{E}_v[\tau] : \tau \text{ is a stopping rule from } v \text{ to } \pi\}$$

and

$$t_{\mathrm{hit}}(\delta, X) = \max_{\substack{v \in V, A \subset V: \\ \pi(A) \geq \delta}} \mathbb{E}_v[H(A)]$$

where $H(A)$ is the first hitting time of A, $H(A) = \min\{n \geq 0 : X_n \in A\}$.

Aldous [1, Theorem 6] showed that there exists a constant c_1 such that for any lazy Markov chain X, and any $\delta > 0$,

$$\delta t_{\mathrm{hit}}(\delta, X) \leq c_1 t_{\mathrm{stop}}(X). \tag{1.1}$$

He also showed [1, Theorem 5] that for lazy *reversible* Markov chains,

$$t_{\mathrm{mix}}(X) \asymp t_{\mathrm{stop}}(X). \tag{1.2}$$

Here the notation $\asymp$ means that there exist constants $0 < c \leq C < \infty$ such that $ct_{\mathrm{mix}}(X) \leq t_{\mathrm{stop}}(X) \leq Ct_{\mathrm{mix}}(X)$. We will use these bounds in our proofs. (In fact Aldous considered continuous time chains, but it is possible to adapt his proofs to discrete time, using the lazy nature of the chain for his Lemma 38.)

Let d_G be the graph distance on V, and let $\Delta(G)$ be the maximum degree over all vertices in G. For $A \subset V$ and $r \geq 0$ write $B_G(r, A) = \{v \in V : d_G(v, A) \leq r\}$ for the closed ball of radius r about A. A *correspondence* between graphs $G = (V, E)$ and $G' = (V', E')$ is a relation $\mathcal{C} \subset V \times V'$ such that for all $v \in V$, there is $v' \in V'$ such that $(v, v') \in \mathcal{C}$; and for all $v' \in V'$, there is $v \in V$ such that $(v, v') \in \mathcal{C}$. In other words, the bipartite graph with vertices $V \cup V'$ and edges $\mathcal{C}$ has no isolated vertices. The *stretch* of $\mathcal{C}$ is

$$\mathrm{str}(\mathcal{C}) = \sup\left\{ \frac{d_{G'}(v', w') \vee d_G(v, w) + 1}{d_{G'}(v', w') \wedge d_G(v, w) + 1} : (v, v') \in \mathcal{C}, (w, w') \in \mathcal{C} \right\}.$$

We say G and G' are *r-roughly isometric*, and write $G \simeq_r G'$, if there is a correspondence $\mathcal{C}$ between G and G' with $\mathrm{str}(\mathcal{C}) \leq r$. Our definition of rough isometry differs slightly from that given in e.g. [3], but the reader may easily check that the two are equivalent up to adjusting the value of r. We prefer the current definition as it is obviously symmetric and will be easier to apply in our setting.

1.2. **Robustness of mixing.**

Theorem 1.1. *Fix a finite irreducible lazy Markov chain X on a graph G, and another such chain, Y, on a tree T. Suppose that X and Y are both ε-uniform and let $\Delta = \Delta(G) \vee \Delta(T)$. If $G \simeq_r T$ then*

$$t_{\mathrm{stop}}(X) \asymp_{\Delta,\varepsilon,r} t_{\mathrm{stop}}(Y).$$

If X and Y are reversible then

$$t_{\mathrm{mix}}(X) \asymp_{\Delta,\varepsilon,r} t_{\mathrm{mix}}(Y).$$

The notation $\asymp_{\Delta,\varepsilon,r}$ indicates that e.g. $(t_{\mathrm{mix}}(X) \vee t_{\mathrm{mix}}(Y))/(t_{\mathrm{mix}}(X) \wedge t_{\mathrm{mix}}(Y))$ is bounded by a function of Δ, ε and r.

Example 1 in Section 2 shows that some dependence on Δ and ε is indeed necessary; obviously dependence on r is also necessary. Furthermore, the aforementioned example of Ding and Peres [5], which we reproduce in Section 2 as Example 3, shows that for fixed Δ and ε the mixing time is not robust under r-rough isometry for any r, so the condition that T must be a tree cannot be removed entirely.

1.3. **Bottleneck sequences.** Our strategy for proving Theorem 1.1 is to give a geometric characterization of the mixing time for treelike graphs. In order to do this we need some further definitions. For $A, B \subset V$ define

$$Q(A,B) = \mathbb{P}_\pi(X_0 \in A, X_1 \in B) \quad \text{and} \quad \Phi(A) = \frac{Q(A,A^c)}{\pi(A)\pi(A^c)}.$$

These quantities capture how easy or difficult it is for X to move between different subsets of the state space. If $A = \{a\}$ is a singleton we write $Q(a,B)$ instead of $Q(\{a\},B)$, and likewise write $Q(A,b)$ and $Q(a,b)$. In the remainder of the paper, we use this convention without comment when applying other set functions to singletons.

Note that for any Markov chain and any $A \subset V$, $Q(V,A) = \mathbb{P}_\pi(X_1 \in A) = \mathbb{P}_\pi(X_0 \in A) = Q(A,V)$ and therefore

$$Q(A,A^c) = Q(A,V) - Q(A,A) = Q(V,A) - Q(A,A) = Q(A^c,A). \tag{1.3}$$

For a set $A \subset V$ of vertices, we write "A is connected" to mean that the induced subgraph $G[A]$ is connected. To clarify, connected means that for any $u, v \in A$ there is a path $u_0, \ldots, u_l$ from u to v such that for each i, either $p_{u_i u_{i+1}} > 0$ or $p_{u_{i+1} u_i} > 0$.

We define $\partial A = \{u \in A^c : \exists v \in A \text{ with } p_{uv} > 0\}$, and $\partial^i A = \{v \in A : \exists u \in A^c \text{ with } p_{uv} > 0\}$.

Given a Markov chain X and $\theta \in (0,1]$, a *θ-bottleneck sequence* for X is an increasing sequence $S_1 \subset S_2 \subset \ldots \subset S_l$ of subsets of V with $S_1 \neq \emptyset$ and $S_l \neq V$ such that

- S_j and S_j^c are both connected for each $j = 1, \ldots, l$;
- $Q(S_{j+1} \setminus S_j, S_j) \geq \theta Q(S_j^c, S_j)$.

The second condition says that, in stationarity, when a random walk enters S_j it is reasonably likely to have come from $S_{j+1} \setminus S_j$. If $\theta = 1$ then it states that $\partial S_j \subset S_{j+1}$. For any $\theta > 0$ it implies that $\partial S_j \cap S_{j+1}$ is non-empty.

Let $\mathcal{S}_\theta = \mathcal{S}_\theta(X)$ be the set of θ-bottleneck sequences for X.

Our proofs are inspired by the approach of Lovasz and Kannan [12], who proved bounds on $t_{\mathrm{mix}}(X)$ for reversible chains by considering bottlenecks at multiple scales, and of Fountoulakis and Reed [6], who showed that connectivity could be exploited to strengthen the Lovasz–Kannan bounds. The next result further improves the bound in [6] by considering only nested sequences of bottlenecks, rather than all possible bottlenecks at each scale. We emphasise that it does not require the underlying graph to be tree-like.

Theorem 1.2. *For any finite irreducible lazy Markov chain X, and any $\theta \in (0,1)$,*

$$t_{\mathrm{stop}}(X) \lesssim_\theta \max_{(S_1,\dots,S_l)\in\mathcal{S}_\theta(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_j)}.$$

If X is also reversible, then

$$t_{\mathrm{mix}}(X) \lesssim_\theta \max_{(S_1,\dots,S_l)\in\mathcal{S}_\theta(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_j)}.$$

The notation $\lesssim_\theta$ means that e.g. $t_{\mathrm{stop}}(X)/\max_{(S_1,\dots,S_l)\in\mathcal{S}_\theta(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_j)}$ is bounded from above by a function of θ only, uniformly in X.

In Section 1.5 we describe the relation between Theorem 1.2 and existing results. In particular, we explain the Fountoulakis-Reed result, and how it follows from Theorem 1.2.

In order to prove Theorem 1.1, we strengthen Theorem 1.2 by showing that we need not maximize over all bottleneck sequences. We again bound $t_{\mathrm{mix}}(X)$ by an expression of the form $\sum 1/\Phi(D_j)$, for some sequence of bottlenecks (D_j). However, the sequence is chosen according to a *game*, rather than simply taking the worst possible sequence as in Theorem 1.2. Informally, this allows us to choose some of the points near D_j and force D_{j+1} to contain those points. This means that it is possible to—roughly speaking—force our bottleneck sequence to move in a particular direction.

1.4. **The bottleneck sequence game.** Fix a finite irreducible lazy Markov chain X as above. We describe a game played between two players, which builds an increasing sequence $D_1,\dots,D_l$ of subsets of V. One player, called Crawler, aims to maximise $\sum_{k=1}^{l} 1/\Phi(D_k)$, usually by making the sequence advance slowly; the other player, Dasher, aims to minimize the same quantity by growing the sequence quickly.

We prove that the rules of the game imply that whatever strategy Dasher adopts, Crawler can make $\sum_{k=1}^{l} 1/\Phi(D_k)$ larger than $t_{\mathrm{stop}}(X)$, up to a constant factor. This is formalized in Theorem 1.3. Knowing this, we can then bound $t_{\mathrm{mix}}(G)$ from above by choosing a specific strategy for Dasher, then proving upper bounds on $\sum_{k=1}^{l} 1/\Phi(D_k)$ when Dasher follows this strategy. This is how we will prove Theorem 1.1.

Defining the game requires a few more definitions. For a set $A \subset V$, recall that $H(A) = \min\{n : X_n \in A\}$ is the first hitting time of A by X, and write $H^+(A) = \min\{n \geq 1 : X_n \in A\}$.

Definition. For $u, v \in V$ and $\alpha \in (0,1]$, we say that v is α-near to u if

$$\pi(v)\mathbb{P}_v(H(u) \le 1/\alpha) \ge \alpha\pi(u).$$

For $A, B \subset V$, we say that A is α-near to B if every vertex of A is α-near to some vertex of B.

Note that if $v \in A$ then clearly v is α-near to A for every α.

Definition. Fix sets $A \subset B \subset V$ with A connected, and $\beta \in (0,1]$. We say that B is a *β-adjustment* of A if for any set $S \subset B^c$ such that $A \cup S$ is connected, it holds that

$$Q((B \cup S)^c, A \cup S) \ge \beta Q((B \cup S)^c, B \cup S)\,.$$

Heuristically, if B is not much bigger than A, and removing $B \setminus A$ from the graph does not "create low-conductance sets containing A", then B is an adjustment of A.

If B is a β-adjustment of A, taking $S = \emptyset$ shows that $Q(B^c, A) \ge \beta Q(B^c, B)$. This states that in stationarity, when the random walk enters B, it is reasonably likely to enter A at the same moment. However, this condition is not sufficient to imply that B is a β-adjustment. To see an example, consider simple random walk on the path $\{1, \dots, n\}$ for $n \ge 4$. Let $A = \{2\}$ and $B = \{2, 3\}$. Then $Q(B^c, A) = Q(B^c, B)/2$. However, if we take $S = \{1\}$, then B separates $A \cup S$ from the rest of the graph so $Q((B \cup S)^c, A \cup S) = 0 < Q((B \cup S)^c, B \cup S)$ and we see that B is not a β-adjustment of A for any $\beta \in (0,1]$.

Fix a vertex $s \in V$, and $\alpha, \beta \in (0,1]$ such that $1/\alpha \in \mathbb{N}$, and $\gamma \in (0,1)$. We now describe the rules of the $(s, \alpha, \beta, \gamma)$-bottleneck sequence game for X. Recall that the players are called Crawler and Dasher. A position of the game is a pair (C, D) of subsets of V, or equivalently an element of $2^V \times 2^V$. We start from $(\emptyset, \emptyset)$. Play alternates, starting with Crawler.

From position (C, D), a γ-valid move for Crawler is any position (C', D) satisfying

(a) **(Connectivity)** $C \subset C'$, $C' \setminus C \subset D^c$, and C' is connected.
(b) **(Isoperimetry)** $Q((D \cup C')^c, C) \le \gamma Q(D^c, C)$.

From position (C, D), an (s, α, β)-valid move for Dasher is any position (C, D') satisfying

(i) **(Complement connectivity)** $D \cup C \subset D'$ and $(D')^c$ is connected.
(ii) **(Nearness)** $\partial^i D'$ is α-near to C.
(iii) **(Adjustment)** D' is a β-adjustment of C.
(iv) **(Endgame)** If $s \in D'$, then s is α-near to C and $D' = V(G)$.

The game ends as soon as Dasher chooses $D' = V(G)$.

Say a sequence $((C_i, D_i), 0 \le i \le \ell + 1)$ is $(s, \alpha, \beta, \gamma)$-valid if $(C_0, D_0) = (\emptyset, \emptyset)$ and

- for each $i \le \ell$, (C_{i+1}, D_i) is a γ-valid move for Crawler starting from (C_i, D_i),
- for each $i \le \ell$, (C_{i+1}, D_{i+1}) is a (α, β, s)-valid move for Dasher starting from (C_{i+1}, D_i).

We say the sequence is an (s,α,β,γ)-valid *game* if it is (s,α,β,γ)-valid and additionally $D_{\ell+1}=V$.

Note that if $((C_i,D_i),0\le i\le \ell+1)$ is a valid game then by isoperimetry, $C_{i+1}\setminus D_i$ is non-empty for each $i\le\ell$. In particular, $D_\ell\neq V$.

In Lemma 6.1 we shall see that the rules are consistent, in that both players can always make a valid move (given that the previous moves have been valid). We now state the upper bound on the mixing time that arises from the bottleneck sequence game, which we will prove in Section 6. Again we emphasise that this bound holds for any graph, not just tree-like graphs.

Theorem 1.3. *Fix a finite irreducible lazy Markov chain X, and any $\alpha,\beta\in(0,1]$ and $\gamma\in(0,1)$. There exist $s\in V$ and γ-valid moves for Crawler such that, whatever (s,α,β)-valid moves Dasher makes,*

$$t_{\mathrm{stop}}(X)\lesssim\frac{1}{\alpha^3}+\frac{1}{\alpha^2\beta\gamma}\sum_{k=1}^{l}\frac{1}{\Phi(D_k)}.$$

If X is also reversible, then

$$t_{\mathrm{mix}}(X)\lesssim\frac{1}{\alpha^3}+\frac{1}{\alpha^2\beta\gamma}\sum_{k=1}^{l}\frac{1}{\Phi(D_k)}.$$

We now briefly explain how to see that Theorem 1.3 implies Theorem 1.2. Given $A,B,C\subset V$, write $A\overset{C}{\longleftrightarrow}B$ if every path from A to B contains a vertex of C. Note that if $a\in C$ then $a\overset{C}{\longleftrightarrow}B$ for all $B\subset V$.

Definition. The *s-hull* of A is the set $h_s(A)=\{t : s\overset{A}{\longleftrightarrow}t\}\supset A$ of vertices disconnected from s when A is removed.

We shall see in Lemma 6.1 that in the course of the game, if the position is (C,D) and it is Dasher's move, then Dasher can always move to $(C,h_s(C\cup D))$. If Dasher follows this strategy, then $D_1,D_2,\dots,D_l$ is a $(1-\gamma)$-bottleneck sequence. Therefore Theorem 1.3 implies that there exist valid moves $C_1,\dots,C_l$ for Crawler such that, if Dasher always moves to $D_k=h_s(C_k\cup D_{k-1})$, then

$$t_{\mathrm{mix}}(X)\lesssim 1+\frac{1}{\gamma}\sum_{k=1}^{l}\frac{1}{\Phi(D_k)}\lesssim\frac{1}{\gamma}\max_{(S_1,\dots,S_l)\in\mathcal{S}_\gamma(X)}\sum_{j=1}^{l}\frac{1}{\Phi(S_j)}.$$

Theorem 1.3 is therefore, up to constants, at least as strong as Theorem 1.2. In some cases it is stronger: Example 4 in Section 2 gives a graph where Theorem 1.2 yields an upper bound on the mixing time of order n^4, whereas Theorem 1.3 gives n^3, which is the correct order.

In Section 7 we use Theorem 1.3 to prove Theorem 1.1.

1.5. **Related work.** Ding, Lee and Peres [4] showed that the expected *cover time* of reversible chains is determined, up to constant factors, by a functional of the graph geometry. Write τ_{cov} for the first time every vertex of V has been visited by X. The cover time for X is defined as $t_{\mathrm{cov}}(X)=\max_{v\in V}\mathbb{E}_v(\tau_{\mathrm{cov}})$.

Fix any $v_0\in V$ and let $\{\eta_v\}_{v\in V}$ be the Gaussian free field (GFF) on G with $\eta_{v_0}=0$. This is the centered Gaussian process whose covariance matrix is determined

by the identities $\mathbb{E}[(\eta_u - \eta_v)^2] = R_{\text{eff}}(u, v)$ for $u, v \in V$. Here, $R_{\text{eff}}(u, v)$ is the effective resistance between u and v when G is viewed as an electrical network with conductances $(c_{uv},\ uv \in E)$ (recall that $c_{uv} = \pi(u)p_{uv}$). Ding, Lee and Peres showed that

$$t_{\text{cov}}(X) \asymp \mathbb{E}\big[\max_{v\in V} \eta_v\big]^2.$$

The maximum of the GFF is a geometric functional, so we may view the cover time as "geometrically robust".

Returning to mixing times, Peres and Sousi [16] showed that for reversible chains, not only does (1.1) hold, but in fact $t_{\text{mix}}(X) \asymp_\delta t_{\text{hit}}(\delta, X)$ for any $\delta < 1/2$. They then used this equivalence to show that if X and Y are two lazy reversible Markov chains on the same tree T with conductances bounded above and below by some strictly positive constants, then

$$t_{\text{mix}}(X) \asymp t_{\text{mix}}(Y). \tag{1.4}$$

Our Theorem 1.2 yields another simple proof of (1.4): see Corollary 4.5. The equivalence $t_{\text{mix}}(X) \asymp_\delta t_{\text{hit}}(\delta, X)$ in the more delicate case $\delta = 1/2$ was later established by Griffiths et al. [8].

Fountoulakis and Reed [6], building on work by Jerrum and Sinclair [10] and Lovász and Kannan [12], showed that if X is reversible and we set

$$\Phi(p) = \min_{\substack{A \text{ connected},\\ p/2 \le \pi(A) \le p}} \Phi(A) \qquad \text{and} \qquad m = \max_{v\in V} \lfloor \log_2 \pi(v)^{-1} \rfloor,$$

then

$$t_{\text{mix}} \lesssim \sum_{j=1}^{m} \frac{1}{\Phi(2^{-j})^2}. \tag{1.5}$$

Our proofs are largely based on the approach of Fountoulakis and Reed [6]. In fact, by following the end of the proof in [6], we may deduce their bound (1.5) from our Theorem 1.2. Example 1 in Section 2 gives a Markov chain for which Theorem 1.2 gives the correct order of magnitude of t_{mix}, whereas (1.5) does not.

Another variant on the Lovász and Kannan result was obtained by Morris and Peres [15], who sharpened (1.5) in several ways, including allowing non-reversible chains and bounding the larger L^∞-mixing time. This mantle was then taken up by Goel, Montenegro and Tetali [7], who used the spectral profile $\Lambda(p)$ instead of the conductance profile $\Phi(p)$. For reversible chains $\Lambda(p)$ may be thought of as the smallest eigenvalue of the Laplacian on the graph restricted to A, minimized over all sets $A \subset V$ of invariant measure at most p; in fact the definition of $\Lambda(p)$ used in [7] is different, but for reversible chains it is equivalent to this definition up to constants provided that $p \le 1/2$. Goel, Montenegro and Tetali gave an upper bound involving $\Lambda(p)$, and were able to recover the result of Morris and Peres [15] in the reversible case by using a discrete Cheeger inequality to relate $\Lambda(p)$ to $\Phi(p)$.

Kozma [11] showed that the upper bound on the L^∞-mixing time given in [7] is not always correct up to constant factors. He then asked the general question "is the mixing time a geometric property?" and conjectured that the mixing time was robust (up to constant factors) under rough isometry for bounded degree graphs.

However, a construction of Ding and Peres [5] shows that (1.4) cannot hold in general if the underlying graph is not a tree, even if it has bounded degree. The

message that we take from this is that the total variation mixing time is *not* geometrically robust in general. One of the main aims of this article is to show that there is robustness amongst a wider class than just trees: indeed, (1.4) holds if the graph has bounded degree and is roughly isometric to a tree. We will explain Ding and Peres' construction in more detail in Section 2.

Hermon [9] has recently shown that the L^∞-mixing time is also not geometrically robust over all bounded degree graphs. His construction combines aspects of Kozma's and Ding and Peres' examples.

1.6. **Layout of the article.** In Section 2 we give several examples of Markov chains that highlight the key features of our results, as well as their limitations. We then begin our proofs, starting in Section 3 with an easy lower bound on the mixing time for Markov chains on trees, which is essentially already known but which gives a framework for proving a similar lower bound on graphs that are roughly isometric to trees. In Section 4 we prove Theorem 1.2, which both serves as a warm-up for the proof of Theorem 1.3 and combines with the work in Section 3 to show robustness for the mixing time on trees (see Corollary 4.5). In Section 5 we generalise the approach in Section 3 to give a lower bound for the mixing time on graphs that are roughly isometric to trees, and then in Section 6 we prove Theorem 1.3. Finally we combine our work from Sections 5 and 6 to prove Theorem 1.1 in Section 7.

2. Examples

In this section we construct some illuminating examples, several of which are referred to elsewhere in the article. We begin with two trees that are roughly isometric, but which have very different mixing times. This example shows that Theorem 1.1 cannot hold without control over the maximum degree of both G and T. It also provides an example where Theorem 1.2 gives the correct order of magnitude for t_{mix}, but (1.5) does not.

Example 1. Take two star graphs of order n (that is, trees with one internal node and n leaves). Choose one leaf from each, v_1 and v_2, and join them by a single edge. Put unit conductance on each edge. This tree T has mixing time of order n: started from any vertex other than v_1 or v_2, it takes time of order n to reach the other half of the graph. However, T is 2-roughly isometric to the tree T' consisting of two nodes joined by a single edge with unit conductance, which obviously has mixing time of order 1.

It is also easy to check that the bound (1.5) gives $t_{\text{mix}}(T) \lesssim n^2$, whereas Theorem 1.2 gives $t_{\text{mix}}(T) \lesssim n$.

Similarly, by considering a path of length 3 whose middle edge has conductance 1 and whose outer edges both have conductance n, which is again roughly isometric to T', we see that Theorem 1.1 cannot hold without some control over the conductance ratio ε.

The next example shows that the bound in Theorem 1.2 need not hold when $\theta = 1$.

Example 2. Take a complete graph on n vertices, and choose one distinguished Hamiltonian cycle. If uv is in the distinguished cycle, then set $c_{uv} = 1$, and

otherwise set $c_{uv} = 1/n^3$. Then the lazy random walk with these conductances takes time of order n^2 to mix, but $\max_{(S_1,\dots,S_l)\in\mathcal{S}_1(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_j)}$ is of order 1.

The next example is due to Ding and Peres [5]. It shows that Theorem 1.1 cannot hold for general bounded degree graphs without some structural assumption, as well as showing that $\max_{(S_1,\dots,S_l)\in\mathcal{S}_1(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_j)}$ is not always a lower bound for t_{mix}. We give a non-rigorous discussion and refer to [5] for the details.

Example 3 (Ding, Peres [5])**.** Take a binary tree T of height K rooted at o. Distinguish the two children of each vertex as the *left* and *right* children. Let L and R be the sets of left and right children respectively. Let $\Gamma_{u,v}$ be the unique path in T from u to v. Define

$$B = \left\{v \in T : K/4 \le |\Gamma_{o,v}| \le K/2, \text{ and } \left||\Gamma_{o,v}\cap L| - |\Gamma_{o,v}\cap R|\right| \le \sqrt{K}\right\}$$

to be the set of *balanced* vertices. To every vertex in B, attach a path of length K. Finally, attach an expander of size $K^2 2^K$ to the leaves of T, in such a way that every leaf is joined to a different vertex in the expander.

It is easy to see that the mixing time is of the same order as the maximum hitting time of the expander over all starting vertices. If every edge has unit conductance, then starting from the root, with high probability we hit of order K balanced vertices, and spend time of order K in each of the attached paths. Therefore the mixing time is at least K^2. On the other hand, if we change the conductance on every edge (u, v) with $v \in L$ to 2, then with high probability we hit at most of order $\sqrt{K}$ balanced vertices before hitting the expander, regardless of where we start. Thus the mixing time is at most $K^{3/2}$.

Note also that taking S_j to be the first j levels of T along with any attached paths for $j = 1, \dots, K-1$, and setting $l = K-1$, we have

$$\sum_{j=1}^{l} \frac{1}{\Phi(S_j)} \gtrsim \sum_{j=K/4}^{K/2} \frac{1}{\Phi(S_j)} \asymp \sum_{j=K/4}^{K/2} K \asymp K^2,$$

provided the conductances all fall between two positive constants. As we have just seen, this is not a lower bound for t_{mix}.

Finally we construct an example where Theorem 1.3 gives the correct order for t_{mix} and Theorem 1.2 does not.

The Cartesian product $G \times H$ of graphs G and H is the graph with vertices $V(G) \times V(H)$ and edges $\{(u,i)(v,j) : \text{either } u = v \text{ and } ij \in E(H), \text{ or } i = j \text{ and } uv \in E(G)\}$.

Example 4. Let $Q = K_k \times C_n$, where K_k is a complete graph with k vertices and C_n is a cycle with n vertices; for concreteness say C_n has edges $c_1c_2, c_2c_3, \dots, c_nc_1$. We say a node of Q has level ℓ if it is an element of the k-clique corresponding to c_ℓ. For this example we think of k as small and n large.

Let $Q_1, \dots, Q_n$ be disjoint copies of Q, and for each $i \le n$ fix a node v_i at level 1 of Q_i. Then create a graph G by adding an edge between v_i and v_{i+1} for each $i \le n-1$. The graph G has kn^2 vertices and maximum degree $k+4$.

Now define a sequence $(S_\ell,\ \ell < n^2)$ as follows. For $i, j \le n$ let $S_{j+n(i-1)}$ contain all nodes of $Q_1, \dots, Q_{i-1}$, and all nodes of Q_i whose level is at most j. When ℓ is a

multiple of n the only edge exiting S_ℓ leads to $S_{\ell+1}$. When ℓ is *not* a multiple of n, the set S_ℓ has $2k+1$ edges to S_ℓ^c, and all but one of these leads to $S_{\ell+1}$. It follows that with unit conductance on every edge, $(S_\ell,\, \ell = 1, \dots, n^2-1)$ is a θ-bottleneck sequence for any $\theta \le 2k/(2k+1)$. Moreover, when $\ell \le n^2/2$ we have $\Phi(S_\ell) \le 2/\ell$, so

$$\sum_\ell \frac{1}{\Phi(S_\ell)} \gtrsim n^4 .$$

However, in reality, the random walk spends on average time $k^2 n$ wandering around each copy of Q each time it visits it, and apart from these detours it traverses a path of length n; so its mixing time is of the order $k^2 n \times n^2 = k^2 n^3$, much smaller than n^4 when $k^2 \ll n$.

Indeed, in the bottleneck sequence game, as soon as Crawler uses vertex v_i in C, Dasher can add the whole of the rest of Q_i to D. In this way the game consists of at most n moves with $\Phi(D_\ell) \ge 1/(2\ell k^2 n)$ for each ℓ; so Theorem 1.3 gives an upper bound on t_{mix} of the order $k^2 n^3$, which agrees with the actual mixing time.

3. An easy lower bound for the mixing time on trees

If our underlying graph is a tree, then we may use the recursive structure to give a simple lower bound for the mixing time, which is essentially a consequence of Moon's theorem [14, Theorem 4.1].

Proposition 3.1. *Suppose that X is a Markov chain on a tree $T = (V, E)$. Then*

$$t_{\mathrm{mix}}(X) \gtrsim \max_{(S_1,\dots,S_l)\in\mathcal{S}_1(X)} \sum_{i=1}^{l} \frac{1}{\Phi(S_i)}.$$

Proof. We note first that every Markov chain on a tree is reversible (indeed, removing any edge uv splits the tree into two connected components A and A^c, so using (1.3), $\pi(u)p_{uv} = Q(A, A^c) = Q(A^c, A) = \pi(v)p_{vu}$) so we may write $c_{uv} = \pi(u)p_{uv} = c_{vu}$ for the conductance of edge uv, and $c_u = \sum_{v\in V} c_{uv} = \pi(u)$.

Suppose $A \subset V$ is such that both A and A^c are connected. Since T is a tree, there is exactly one edge uv with $u \in A$ and $v \in A^c$.

The observation that $\mathbb{E}_v[H^+(v) \,|\, X_1 = u] = \mathbb{E}_u[H(v)] + 1$, and the identity (see [2, Lemma 2.5]) $\mathbb{E}_v[H^+(v)] = \pi(v)^{-1}$ applied in the subtree of T with vertices $A \cup \{v\}$, together yield that

$$\mathbb{E}_u[H(v)] + 1 = \frac{c_{uv} + \sum_{x\in A} c_x}{c_{uv}} = \frac{Q(A, A^c) + \pi(A)}{Q(A^c, A)} = 1 + \frac{\pi(A)}{Q(A, A^c)}.$$

Thus $\mathbb{E}_u[H(v)] \ge 1/\Phi(A)$.

Choose a 1-bottleneck sequence $S_1, \dots, S_l$ for X. Let $L = \max\{i : \pi(S_i) \le 1/2\}$. Since each set S_i is connected and has connected complement, for each $i \le l$ there is a unique edge $s_i t_i$ with $s_i \in S_i$ and $t_i \in S_i^c$. If the Markov chain starts in S_1, then in order to visit S_L^c the chain must cross each of the edges $s_i t_i$ for $i \le L$. It follows by (1.1) and (1.2) and the above that for any $v \in S_1$,

$$t_{\mathrm{mix}} \gtrsim \mathbb{E}_v[H(S_L^c)] \ge \sum_{i=1}^{L} \mathbb{E}_{s_i}[H(t_i)] \ge \sum_{i=1}^{L} \frac{1}{\Phi(S_i)} .$$

Finally, note that the sequence $(S_l^c, S_{l-1}^c, \ldots, S_1^c)$ is also a 1-bottleneck sequence; applying the above argument to this sequence gives $t_{\text{mix}} \gtrsim \sum_{i=L+1}^{l} 1/\Phi(S_i)$, and the result follows. □

4. Proof of Theorem 1.2

Although Theorem 1.2 is essentially a special case of Theorem 1.3, its proof contains several ideas (and two lemmas and a corollary) that we will need later, so we include it as a warm-up. We recall that for $s \in V$, a stopping rule τ from s to π is simply a stopping time that may use additional randomness to decide when to stop, and satisfies $\mathbb{P}_s(X_\tau = v) = \pi(v)$ for all $v \in V$. It is known (see [13]) that *optimal* stopping rules always exist, i.e. there is a stopping rule τ satisfying $\mathbb{E}_s[\tau] = \inf\{\mathbb{E}_s[\tau'] : \tau' \text{ is a stopping rule from } s \text{ to } \pi\}$. By (1.2), Theorem 1.2 is therefore a consequence of the following result.

Proposition 4.1. *Suppose that X is a Markov chain on a graph $G = (V, E)$ with invariant measure π. Fix $\theta \in (0, 1)$. For any $s \in V$ and any optimal stopping rule τ from s to π,*

$$\mathbb{E}_s[\tau] < \frac{1}{1-\theta} \max_{(S_1,\ldots,S_l)\in\mathcal{S}_\theta(X)} \sum_{j=1}^{l} \frac{1}{\Phi(S_k)}.$$

Our aim now is to prove Proposition 4.1. Our strategy is similar to that of Fountoulakis and Reed [6].

Fix $s \in V$ and an optimal stopping rule τ from s to π. For $v \in V$, define the *scaled exit frequencies* $\{y_v, v \in V\}$ by

$$y_v = \frac{1}{\pi(v)} \mathbb{E}_s[\#\{k < \tau : X_k = v\}].$$

(Since we will be keeping s and τ fixed, we omit them from the notation.) Label the elements of V as $1, \ldots, N$ so that $y_1 \le \ldots \le y_N$. We will bound $\mathbb{E}_s[\tau]$ using the identity $\mathbb{E}_s[\tau] = \sum_{v=1}^{N} \pi(v) y_v$.

Since τ is an optimal stopping rule, it has a *halting state* h such that $y_h = 0$ (again see [13]). Thus $y_h = y_1 = 0$ and we may certainly choose the ordering above so that $h = 1$. The following identity is key to our analysis. It is originally due to Lovasz and Kannan [12], but we include a short proof as our formulation is slightly different; in particular we allow non-reversible chains.

Lemma 4.2. *If $Z \subset V$ and $s \notin Z$ then*

$$\pi(Z) = \sum_{u\notin Z, v\in Z} (y_u Q(u, v) - y_v Q(v, u)).$$

Proof. For all $v \in V$, $X_\tau = v$ if and only if $\#\{k \le \tau : X_k = v\} = \#\{k < \tau : X_k = v\} + 1$, so

$$\pi(v) = \mathbb{P}_s(X_\tau = v) = \mathbb{E}_s[\#\{k \le \tau : X_k = v\}] - \mathbb{E}_s[\#\{k < \tau : X_k = v\}].$$

If $v \neq s$ then by the Markov property,

$$\begin{aligned}\mathbb{E}_s[\#\{k \leq \tau : X_k = v\}] &= \sum_u \mathbb{E}_s[\#\{k < \tau : X_k = u, X_{k+1} = v\}] \\ &= \sum_u y_u \pi(u) p_{uv} = \sum_u y_u Q(u,v)\,.\end{aligned}$$

Since

$$\mathbb{E}_s[\#\{k < \tau : X_k = v\}] = y_v \pi(v) = \sum_u y_v \pi(v) p_{vu} = \sum_u y_v Q(v,u)\,,$$

we obtain $\pi(v) = \sum_u (y_u Q(u,v) - y_v Q(v,u))$. For $Z \subset V$ with $s \notin Z$, we thus have

$$\pi(Z) = \sum_{u \in V} \sum_{v \in Z} (y_u Q(u,v) - y_v Q(v,u))$$

and the result follows since, by symmetry, $\sum_{u,v \in Z} (y_u Q(u,v) - y_v Q(v,u)) = 0$. □

Corollary 4.3. *For any $i \in V$, the set $\mathcal{A}(i) := \{j \in V : y_j \geq y_i\}$ is internally connected from s: that is, for every vertex v in $\mathcal{A}(i)$, there exists a path from s to v within $\mathcal{A}(i)$.*

Proof. Let Z be the set of vertices v in $\mathcal{A}(i)$ such that there does not exist a path from s to v within $\mathcal{A}(i)$. For any $k \in Z$ and $\ell \notin Z$, if $Q(\ell,k) > 0$ then necessarily $\ell \notin \mathcal{A}(i)$ (otherwise there would be a path from s to k within $\mathcal{A}(i)$) so we have $y_\ell < y_i$. Also, for any $k \in Z$ we have $y_k \geq y_i$. Applying Lemma 4.2, since $s \notin Z$,

$$\pi(Z) = \sum_{\ell \notin Z} \sum_{k \in Z} (y_\ell Q(\ell,k) - y_k Q(k,\ell)) \leq y_i Q(Z^c, Z) - y_i Q(Z, Z^c)\,,$$

with equality if and only $Z = \emptyset$. But (1.3) tells us that $Q(Z^c, Z) = Q(Z, Z^c)$, so the right-hand side is zero and the result follows. □

It follows that y_s is at least as large as any of the other y_v, so $y_s = y_N$ and we may assume that $s = N$. (Recall that we also have $h = 1$ where h is the halting state, and $y_1 = 0$.)

For $i = 1, \ldots, N$, let B_i be the set of vertices in $\{1, \ldots, i\}$ that are internally connected to 1, i.e. the set of $v \in \{1, \ldots, i\}$ such that there exists a path from v to 1 within $\{1, \ldots, i\}$.

Lemma 4.4. *For all $1 \leq i \leq j < N$,*

$$y_{j+1} - y_i \leq \frac{\pi(B_i)}{Q(B_j^c, B_i)}\,.$$

Proof. Since for $i < N$ we know that $s \notin B_i$, by Lemma 4.2 we have

$$\begin{aligned}\pi(B_i) &= \sum_{\ell \notin B_i} \sum_{k \in B_i} (y_\ell Q(\ell,k) - y_k Q(k,\ell)) \\ &\geq \sum_{\ell \notin B_i} \sum_{k \in B_i} y_\ell Q(\ell,k) - y_i Q(B_i, B_i^c) \\ &= \sum_{\ell \notin B_j} \sum_{k \in B_i} y_\ell Q(\ell,k) + \sum_{\ell \in B_j \setminus B_i} \sum_{k \in B_i} y_\ell Q(\ell,k) - y_i Q(B_i, B_i^c).\end{aligned}$$

If $\ell \in B_j \setminus B_i$ and $Q(\ell, k) > 0$ for some $k \in B_i$ then $y_\ell \geq y_i$, so the preceding bound gives

$$\pi(B_i) \geq \sum_{\ell \notin B_j} \sum_{k \in B_i} y_\ell Q(\ell, k) + y_i Q(B_j \setminus B_i, B_i) - y_i Q(B_i, B_i^c).$$

Since (see (1.3)) $Q(A, A^c) = Q(A^c, A)$ for any A, we get

$$\pi(B_i) \geq \sum_{\ell \notin B_j} \sum_{k \in B_i} y_\ell Q(\ell, k) - y_i Q(B_j^c, B_i).$$

Finally, if $l \notin B_j$, $k \in B_i$, and $Q(l, k) > 0$, then $y_l \geq y_{j+1}$. Therefore

$$\pi(B_i) \geq (y_{j+1} - y_i) Q(B_j^c, B_i)\,. \qquad \square$$

We can now prove our main result for this section.

Proof of Proposition 4.1. Fix $\theta \in (0, 1)$. Let $m_1 = 1$ and for $i \geq 1$, define

$$m_{i+1} = \min\{m > m_i : Q(B_m^c, B_{m_i}) \leq (1 - \theta) Q(B_{m_i}^c, B_{m_i})\}\,,$$

or $m_{i+1} = N$ if no such m exists. Let $l = \min\{i : m_i = N\} - 1$. Note that for each j, by definition B_j is internally connected to 1 within $\{1, \ldots, j\}$. Since our Markov chain is irreducible, for any $i \leq j$ such that $i \notin B_j$, there must exist a path from i to k for some $k > j$. Since $\{j + 1, \ldots, N\}$ is internally connected from N by Corollary 4.3, it follows that B_j^c must be connected. Also, for any $i \leq l$, by the definition of m_{i+1} we have

$$(1 - \theta) Q(B_{m_i}^c, B_{m_i}) \geq Q(B_{m_{i+1}}^c, B_{m_i}) = Q(B_{m_i}^c, B_{m_i}) - Q(B_{m_{i+1}} \setminus B_{m_i}, B_{m_i})$$

and rearranging we get

$$Q(B_{m_{i+1}} \setminus B_{m_i}, B_{m_i}) \geq \theta Q(B_{m_i}^c, B_{m_i});$$

thus $(B_{m_i}, 1 \leq i \leq l)$ is a θ-bottleneck sequence.

Now, since $y_1 = 0$, we have

$$\begin{aligned}
\mathbb{E}_s[\tau] = \sum_{v=1}^{N} \pi(v) y_v &\leq \sum_{v=1}^{N} \pi(v) \cdot \sum_{j=1}^{l} (y_{m_{j+1}} - y_{m_j}) \mathbb{1}_{\{v > m_j\}} \\
&= \sum_{j=1}^{l} (y_{m_{j+1}} - y_{m_j}) \pi(\{m_j + 1, \ldots, N\}) \\
&\leq \sum_{j=1}^{l} (y_{m_{j+1}} - y_{m_j}) \pi(B_{m_j}^c) \\
&\leq \sum_{j=1}^{l} \frac{\pi(B_{m_j}) \pi(B_{m_j}^c)}{Q(B_{m_{j+1}-1}^c, B_{m_j})}\,,
\end{aligned}$$

the last inequality by Lemma 4.4.

By the definition of m_{j+1}, we have $Q(B_{m_{j+1}-1}^c, B_{m_j}) > (1-\theta) Q(B_{m_j}^c, B_{m_j})$, so the preceding inequality yields

$$\mathbb{E}_s[\tau] < \frac{1}{1 - \theta} \sum_{j=1}^{\ell} \frac{1}{\Phi(B_{m_j})}\,. \qquad \square$$

This completes the proof of Theorem 1.2. We now note that if our graph is a tree, then any connected set $A \subset V$ with connected complement has a single edge between A and A^c, and therefore every θ-bottleneck sequence (for $\theta \in (0,1]$) is a 1-bottleneck sequence. We also recall that the same property implies that every Markov chain on a tree is reversible, and we may therefore unambiguously set $c_{uv} = \pi(u)p_{uv}$ for each $uv \in V$. Combining Theorem 1.2 with Proposition 3.1, we obtain:

Corollary 4.5. *Suppose that X is a lazy Markov chain on a tree. Then*

$$t_{\mathrm{mix}}(X) \asymp \max_{(S_1,\dots,S_l)\in\mathcal{S}_1(T)} \sum_{i=1}^{l} \frac{1}{\Phi(S_l)}.$$

Therefore if X and Y are two lazy Markov chains on the same tree, with conductances satisfying $\varepsilon c^Y_{uv} \le c^X_{uv} \le c^Y_{uv}/\varepsilon$ for all $uv \in E$ and some $\varepsilon > 0$, we have

$$t_{\mathrm{mix}}(X) \asymp_\varepsilon t_{\mathrm{mix}}(Y).$$

The latter part of this result was originally proved by Peres and Sousi [16] using very different methods.

5. Lower bounds for the mixing time on more general graphs

Proposition 3.1 gives a lower bound for the mixing time on trees in terms of 1-bottleneck sequences. The proof uses the tree structure in a non-trivial way, but the idea, that in order to mix we must be able to hit every large set, and to hit a large set we must travel through a sequence of bottlenecks, holds more generally. Indeed, if G is roughly isometric to a tree T, then a similar approach can be used to give a lower bound on the mixing time.

We recall some definitions: for any graph G, d_G denotes the graph distance on V. For $A \subset V$ and $r \ge 0$ we write $B_G(r, A) = \{v \in V : d_G(v, A) \le r\}$, and for $A, B, C \subset V$ we write $A \overset{C}{\longleftrightarrow} B$ if every path between A and B contains a vertex of C.

Proposition 5.1. *Fix a finite connected graph $G = (V(G), E(G))$ and a finite tree $T = (V(T), E(T))$. Suppose that X is an ε-uniform Markov chain on G, and Y is an ε-uniform Markov chain on T. Suppose also that $G \simeq_r T$, and G and T both have maximum degree at most Δ. Then there exists $\eta > 0$, depending only on Δ, ε and r, such that*

$$t_{\mathrm{hit}}(\eta, X) \gtrsim_{\Delta,\varepsilon,r} \max_{(S_1,\dots,S_l)\in\mathcal{S}_1(T)} \sum_{i=1}^{l} \frac{|S_i||S_i^c|}{|V(T)|} \asymp_{\Delta,\varepsilon} \max_{(S_1,\dots,S_l)\in\mathcal{S}_1(T)} \sum_{i=1}^{l} \frac{1}{\Phi^Y(S_i)}.$$

To prove this, our strategy is as follows. Take a 1-bottleneck sequence for T. Between any set S_j in that sequence and its complement S_j^c, there is a single edge. The same is not necessarily true for bottleneck sequences on G, but we can use the rough isometry between T and G to build a small subset of $V(G)$ that plays the role of the single edge in T, in that when this subset is removed from G it splits the graph into at least two components. These two components will be of similar sizes to S_j and S_j^c. In order to mix, the Markov chain will have to pass through all of the small subsets, and this will yield the required bound.

In order to make this heuristic precise, fix a correspondence $\mathcal{C} \subset V(T) \times V(G)$ with $\mathrm{str}(\mathcal{C}) \leq r$. For $t \in V(T)$ write $\mathcal{C}_{t,G} = \{v \in V(G) : (t,v) \in \mathcal{C}\}$, and likewise for $v \in V(G)$ let $\mathcal{C}_{T,v} = \{t \in V(T) : (t,v) \in \mathcal{C}\}$.

Next, for $t \in V(T)$ let

$$A_t = \{v \in V(G) : d_G(v, \mathcal{C}_{t,G}) \leq r(r+1)\} = B_G(r(r+1), \mathcal{C}_{t,G}). \tag{5.1}$$

We will need some basic properties of the sets A_t before we can proceed.

Lemma 5.2. *For all $t \in V(T)$, A_t is connected. Moreover, for all $u, v \in A_t$, there is a path P from u to v contained within A_t, of length at most $2r(r+1) + (r-1)$.*

Proof. If $u, v \in \mathcal{C}_{t,G}$ then $d_G(u,v) \leq r-1$ since $\mathrm{str}(\mathcal{C}) \leq r$. It follows that $B_G(\lceil (r-1)/2 \rceil, \mathcal{C}_{t,G})$ is connected, so A_t is also connected. To prove the second claim, note that for any $u, v \in A_t$, $d_G(u,v) \leq d_G(u, \mathcal{C}_{t,G}) + d_G(v, \mathcal{C}_{t,G}) + r - 1$, and that any shortest path between elements of $\mathcal{C}_{t,G}$ is contained within A_t. □

The next lemma relates the components of A_t^c (in G) with those of $V(T) \setminus \{t\}$ (in T). For $x, y \in V(T)$ write $[\![x, y]\!]$ for the unique path between x and y in T.

Lemma 5.3. *For all distinct $a, b, t \in V(T)$, if $t \in [\![a, b]\!]$ then $\mathcal{C}_{a,G} \overset{A_t}{\longleftrightarrow} \mathcal{C}_{b,G}$.*

Proof. Fix distinct $a, b, t \in V(T)$ with $t \in [\![a, b]\!]$, and $x \in \mathcal{C}_{a,G}$ and $y \in \mathcal{C}_{b,G}$, and consider any path $x = x_0, x_1, \ldots, x_\ell = y$ between x and y in G.

For each $i < \ell$, if $a' \in \mathcal{C}_{T,x_i}$ and $b' \in \mathcal{C}_{T,x_{i+1}}$ then since $\mathrm{str}(\mathcal{C}) \leq r$ we have

$$d_T(a', b') + 1 \leq r(d_G(x_i, x_{i+1}) + 1) = 2r.$$

It follows that $\bigcup_{i \leq \ell} B_T(r, \mathcal{C}_{T,x_i})$ contains a path from a to b. Since T is a tree, any such path contains t, so this implies there is $i \leq \ell$ and $s \in \mathcal{C}_{T,x_i}$ with $d_T(s,t) \leq r$.

Finally, fix any vertex $v \in \mathcal{C}_{t,G}$. Then since $(s, x_i) \in \mathcal{C}$ and $(t, v) \in \mathcal{C}$, again using that $\mathrm{str}(\mathcal{C}) \leq r$, we obtain that

$$d_G(x_i, v) + 1 \leq r(d_T(s,t) + 1) \leq r(r+1),$$

so $x_i \in A_t$. Since the path between x and y was arbitrary, it follows that $x \overset{A_t}{\longleftrightarrow} y$. □

Lemma 5.4. *For any $s, t \in V(T)$, if $d_T(s,t) \geq 2r^2(r+1) + r$ then $A_s \cap A_t = \emptyset$.*

Proof. If $x \in \mathcal{C}_{s,G}$ and $y \in \mathcal{C}_{t,G}$ then $d_G(x,y) + 1 \geq (d_T(s,t)+1)/r > 2r(r+1)+1$. The lemma follows. □

A large part of our proof of Proposition 5.1 essentially boils down to invoking Kac's formula. See for example [2, Corollary 2.24] for a standard form of this result; however, we need to apply it in a non-standard way, and so the alternative form below will be useful.

Lemma 5.5 (Kac's formula)**.** *Suppose that X is a finite irreducible Markov chain with state space V. If L, C and R partition V and $L \overset{C}{\longleftrightarrow} R$ then*

$$\pi(L \cup C) = \pi(C) \mathbb{E}_{\pi_C}[H^+(L^c)]$$

where π_C is the probability measure on V given by $\pi_C(v) = (\pi(v)/\pi(C)) \mathbb{1}_{\{v \in C\}}$.

Proof. If $k \geq 2$, then

$$\begin{aligned}
&\mathbb{P}_\pi(H^+(C) = k, X_1 \in L \cup C) \\
&= \mathbb{P}_\pi(X_1 \in L, \dots, X_{k-1} \in L, X_k \in C) \\
&= \mathbb{P}_\pi(X_1 \in L, \dots, X_{k-1} \in L) - \mathbb{P}_\pi(X_1 \in L, \dots, X_k \in L) \\
&= \mathbb{P}_\pi(X_1 \in L, \dots, X_{k-1} \in L) - \mathbb{P}_\pi(X_0 \in L, \dots, X_{k-1} \in L) \\
&= \mathbb{P}_\pi(X_0 \in C, X_1 \in L, \dots, X_{k-1} \in L) \\
&= \pi(C)\mathbb{P}_{\pi_C}(H^+(L^c) \geq k).
\end{aligned}$$

where we used stationarity for the third equality. If $k = 1$, then

$$\mathbb{P}_\pi(H^+(C) = k, X_1 \in L \cup C) = \mathbb{P}_\pi(X_1 \in C) = \pi(C) = \pi(C)\mathbb{P}_{\pi_C}(H^+(L^c) \geq 1).$$

Summing over k, we get

$$\pi(L \cup C) = \mathbb{P}_\pi(X_1 \in L \cup C) = \pi(C)\mathbb{E}_{\pi_C}[H^+(L^c)]. \qquad \square$$

Now fix a 1-bottleneck sequence $(S_1, \dots, S_l)$ for T, and let $M = \max\{i : \pi(S_i) \leq 1/2\}$. For each $i = 1, \dots, M$, let t_i be the unique vertex that is in S_i and has a neighbour in S_i^c. Let $m_1 = M$, and for $j \geq 2$, define

$$m_j = \max\{m : d_T(t_m, t_{m_{j-1}}) \geq 2r^2(r+1) + r\}$$

or $m_j = 0$ if no such m exists. Let $K = \max\{j : m_j \geq 1\}$.

For $j = 1, \dots, K$, let $n_j = m_{K-j+1}$, and define

$$C_j = A_{t_{n_j}}, \quad L_j = \bigcup_{t \in S_{n_j}} \mathcal{C}_{t,G} \setminus A_{t_{n_j}}, \quad R_j = V \setminus (L_j \cup C_j) = \bigcup_{t \in S_{n_j}^c} \mathcal{C}_{t,G} \setminus A_{t_{n_j}}.$$

The next lemma does most of the work in proving our main result for this section, Proposition 5.1.

Lemma 5.6. *If $v \in C_1$, then*

$$\mathbb{E}_v[H(R_K)] \gtrsim_{\Delta,\varepsilon,r} \sum_{i=1}^{M} |S_i|.$$

Proof. By Lemma 5.4, we know that the sets $C_1, \dots, C_K$ are pairwise disjoint. By Lemma 5.3 we know that for each j, $L_j \overset{C_j}{\longleftrightarrow} R_j$. In particular we have $C_{j-1} \subset L_j$ for each $j = 2, \dots, K$, and $C_j \subset R_{j+1}$ for each $j = 1, \dots K-1$.

Therefore if $v \in C_1$,

$$\begin{aligned}
\mathbb{E}_v[H(R_K)] &\geq \mathbb{E}_v[H(R_1)] + \min_{u \in C_2} \mathbb{E}_u[H(R_2)] + \dots + \min_{u \in C_L} \mathbb{E}_u[H(R_K)] \\
&\geq \sum_{j=1}^{K} \min_{u \in C_j} \mathbb{E}_u[H(R_j)].
\end{aligned}$$

Now, if $u \in C_j$, by Lemma 5.2 we have

$$\mathbb{E}_u[H(R_j)] \geq \mathbb{E}_u[H^+(L_j^c)] \asymp_{\Delta,\varepsilon,r} \mathbb{E}_{\pi_{C_j}}[H^+(L_j^c)].$$

By Lemma 5.5, this equals $\pi(L_j \cup C_j)/\pi(C_j)$, and applying Lemma 5.2 again we get

$$\mathbb{E}_v[H(R_K)] \gtrsim_{\Delta,\varepsilon,r} \sum_{j=1}^{K} \frac{\pi(L_j \cup C_j)}{\pi(C_j)} \asymp_{\Delta,\varepsilon,r} \sum_{j=1}^{K} |L_j \cup C_j| \asymp_{\Delta,\varepsilon,r} \sum_{j=1}^{K} |S_{n_j}|.$$

Finally, we have $n_{j+1} \le n_j + 2r^2(r+1) + r$ for each j. We also have $|S_i| \le |S_j|$ for all $i \le j$. Therefore $\sum_{i=1}^{M} |S_i| \le (2r^2(r+1)+r) \sum_{j=1}^{K} |S_{n_j}|$, which completes the proof. □

Proof of Proposition 5.1. Since $R_K = \bigcup_{t \in S_M^c} \mathcal{C}_{t,G} \setminus A_{t_M}$ and we know that $\pi(S_M^c) \ge 1/2$, assuming that G and T are both large (otherwise the result is trivial) by Lemma 5.2 we have $\pi(R_K) \gtrsim_{\Delta,\varepsilon,r} 1$. Thus there exists $\eta > 0$ depending only on Δ, ε and r such that

$$t_{\mathrm{hit}}(\eta, X) \ge \max_{v \in C_1} \mathbb{E}[H(R_K)].$$

Applying Lemma 5.6 we get

$$t_{\mathrm{hit}}(\eta, X) \gtrsim_{\Delta,\varepsilon,r} \sum_{i=1}^{M} |S_i| \ge \sum_{i=1}^{M} \frac{|S_i||S_i^c|}{|V(T)|}.$$

Since $S_l^c, S_{l-1}^c, \dots, S_1^c$ is also a 1-bottleneck sequence, applying the same argument to this sequence gives $t_{\mathrm{hit}}(\eta, X) \gtrsim_{\Delta,\varepsilon,r} \sum_{i=M+1}^{l} \frac{|S_i||S_i^c|}{|V(T)|}$. The result follows. □

6. Proof of Theorem 1.3

We recall the bottleneck sequence game from Section 1.4, and aim to prove Theorem 1.3. First we check that both players can always move.

Lemma 6.1. *Fix $\alpha, \beta \in (0,1]$, $\gamma \in (0,1)$ and $s \in V$. Suppose that $((C_i, D_i), 0 \le i \le k)$ is an $(s, \alpha, \beta, \gamma)$-valid sequence and $D_i \ne V$. Then Crawler has at least one γ-valid move (C_{k+1}, D_k) from position (C_k, D_k). Moreover, from any such position (C_{k+1}, D_k), Dasher has at least one (s, α, β)-valid move (C_{k+1}, D_{k+1}).*

Proof. Fix $((C_i, D_i), 0 \le i \le k)$ as in the statement of the lemma. We begin by showing that Crawler has a valid move. More specifically, writing $C' = C_k \cup D_k^c$, we show that (C', D_k) is a γ-valid move for Crawler starting from (C_k, D_k)

If $k = 0$ then $C_k = D_k = \emptyset$ and $C' = D_k^c = V$, in which case it is easy to see that (C', D_k) is a valid move for Crawler. If $k > 0$ and $D_k \ne V$ then first note that C_k and D_k^c are both connected since $((C_i, D_i), 0 \le i \le k)$ is a valid sequence. Moreover, using the adjustment property for D_k, with $S = 0$, we have that $Q(D_k^c, C_k) \ge \beta Q(D_k^c, D_k) > 0$, so C' is connected. Next, $Q((D_k \cup C')^c, C_k) = Q(\emptyset, C_k) = 0$, so C' satisfies the isoperimetry requirement. Thus (C', D_k) is also a valid move for Crawler in the case $k > 0$.

Now let (C_{k+1}, D_k) be any γ-valid move for Crawler starting from (C_k, D_k), and let $D' = h_s(C_{k+1} \cup D_k)$, the set of all vertices separated from s by $C_{k+1} \cup D_k$. We claim that (C_{k+1}, D') is a (s, α, β)-valid move for Dasher from position (C_{k+1}, D_k).

First, since $D_k \ne V$, the endgame property (iv) applied to (C_k, D_k) yields that $s \notin D_k$. It follows that $s \in D'$ if and only if $s \in C_{k+1}$. In this case $D' =$

$h_s(C_{k+1} \cup D_k) = V(G)$, which verifies (iv). In this case (i) is obvious, and (ii) and (iii) are vacuously true, so (C_{k+1}, D') is a valid move in this case.

We henceforth assume that $s \notin C_{k+1}$ so $s \notin D'$, and (iv) is therefore satisfied. In this case, by definition $(D')^c$ is the connected component of $(C_{k+1} \cup D_k)^c$ containing s. In particular, $(D')^c$ is connected so (i) is satisfied.

In what follows it is convenient to write $R = D' \setminus (C_{k+1} \cup D_k)$, so R consists of all connected components of $(C_{k+1} \cup D_k)^c$ except the one containing s. Note that there are no edges between R and $(D')^c$.

Fix $v \in \partial^i D'$ and a neighbour w of v with $w \in (D')^c$. By the observation of the preceding paragraph, $v \notin R$ so either $v \in C_{k+1}$ or $v \in D_k$. If $v \in C_{k+1}$ then v is certainly α-near to C_{k+1}. If $v \in D_k$ then since $w \in D_k^c$ we have $v \in \partial^i D_k$. By property (ii) applied to (C_k, D_k) we obtain that v is α-near to C_k and thus to C_{k+1}. This establishes (ii).

It remains to show that D' is a β-adjustment of C_{k+1}. For this fix any set $S \subset (D')^c$ for which $C_{k+1} \cup S$ is connected, and let $S' = (D' \setminus D_k) \cup S$.

We wish to apply (iii) to (C_k, D_k) and S', and for this we need that $C_k \cup S'$ is connected. To establish this, first note that $(C_{k+1} \setminus C_k) \subset D_k^c$, so $D' \setminus D_k = (C_{k+1} \setminus C_k) \cup R$. It follows that

$$C_k \cup S' = C_k \cup (C_{k+1} \setminus C_k) \cup R \cup S = C_{k+1} \cup R \cup S\,.$$

Since D_k^c is connected by (i), every component of $(C_{k+1} \cup D_k)^c$ must have an edge between it and C_{k+1}. It follows that $C_{k+1} \cup R$ is connected. But $C_{k+1} \cup S$ is connected by assumption, so $C_k \cup S' = C_{k+1} \cup R \cup S$ is indeed connected.

Applying (iii) to (C_k, D_k) and S', we now obtain that

$$Q((D_k \cup S')^c, C_k \cup S') \geq \beta Q((D_k \cup S')^c, D_k \cup S')\,.$$

Noting that $D_k \cup S' = D' \cup S$ and $C_k \cup S' = C_{k+1} \cup R \cup S$, this inequality becomes

$$Q((D' \cup S)^c, C_{k+1} \cup R \cup S) \geq \beta Q((D' \cup S)^c, D' \cup S)\,.$$

But R consists of all components of $(C_{k+1} \cup D_k)^c$ except the one containing s, and $(D')^c$ is the component of $(C_{k+1} \cup D_k)^c$ containing s, so $Q((D' \cup S)^c, R) \leq Q((D')^c, R) = 0$. We thus have

$$Q((D' \cup S)^c, C_{k+1} \cup S) \geq \beta Q((D' \cup S)^c, D' \cup S)\,,$$

so D' is a β-adjustment of C_{k+1}. This establishes (iii) for D' and completes the proof. □

At this point we recall the scaled exit frequencies $(y_v,\, v \in V)$ from Section 4, and observe the following consequence of nearness.

Lemma 6.2. *Suppose that τ is a stopping rule from s to π. For any $u, v \in V$ and $m \in \mathbb{N}$, if v is $1/m$-near to u then*

$$y_v \leq m(m+1)y_u + m^2(m+1)/2\,.$$

Proof. For any $u, v \in v$ we have

$$\begin{aligned}
&\mathbb{E}_s[\#\{j < \tau : X_j = u\}] \\
&\geq \mathbb{E}_s[\#\{j < \tau - k : X_j = v,\, X_{j+k} = u\}] \\
&= \sum_{j=0}^{\infty} \mathbb{P}_s(X_j = v,\, X_{j+k} = u,\, \tau > j + k) \\
&= \sum_{j=0}^{\infty} \mathbb{P}_s(X_j = v,\, X_{j+k} = u,\, \tau > j) - \sum_{j=0}^{\infty} \mathbb{P}_s(X_j = v,\, X_{j+k} = u,\, j < \tau \leq j + k).
\end{aligned}$$

By the simple Markov property,

$$\mathbb{P}_s(X_j = v,\, X_{j+k} = u,\, \tau > j) = \mathbb{P}_s(X_j = v\, \tau > j)\mathbb{P}_v(X_k = u),$$

so

$$\begin{aligned}
&\mathbb{E}_s[\#\{j < \tau : X_j = u\}] \\
&\geq \sum_{j=0}^{\infty} \mathbb{P}_s(X_j = v\, \tau > j)\mathbb{P}_v(X_k = u) - \sum_{j=0}^{\infty} \mathbb{P}_s(X_{j+k} = u,\, j < \tau \leq j + k) \\
&\geq \mathbb{E}_s[\#\{j < \tau : X_j = v\}]\mathbb{P}_v(X_k = u) - \mathbb{E}_s[\#\{i \in \{\tau, \ldots, \tau + k - 1\} : X_i = u\}] \\
&= \mathbb{E}_s[\#\{j < \tau : X_j = v\}]\mathbb{P}_v(X_k = u) - \sum_{j=0}^{k-1} \mathbb{P}_s(X_{\tau+j} = u).
\end{aligned}$$

But $X_{\tau+j}$ has distribution π for all $j \geq 0$, so recalling that $\mathbb{E}_s[\#\{j < \tau : X_j = u\}] = \pi(u)y_u$, and similarly for v, we get

$$\pi(u)y_u \geq \pi(v)y_v\mathbb{P}_v(X_k = u) - k\pi(u).$$

Finally, for any $m \in \mathbb{N}$ we have $\mathbb{P}_v(H(u) \leq m) \leq \sum_{k=0}^{m} \mathbb{P}_v(X_k = u)$, so summing the preceding bound over $k \leq m$ gives

$$(m + 1)\pi(u)y_u \geq \pi(v)y_v\mathbb{P}_v(H(u) \leq m) - m(m + 1)\pi(u)/2.$$

Dividing through by $\pi(u)$, we have

$$(m + 1)y_u \geq \frac{\pi(v)y_v}{\pi(u)} \cdot \mathbb{P}_v(H(u) \leq m) - m(m + 1)/2.$$

If v is $(1/m)$-near to u we have $\pi(v)\mathbb{P}_v(H(u) \leq m) \geq \pi(u)/m$, and the result follows. □

In order to prove Theorem 1.3, by (1.2), it is enough to prove the following.

Proposition 6.3. *Fix a graph $G = (V, E)$, and suppose that X is a Markov chain on G. Take any $s \in V$, $\alpha, \beta \in (0, 1]$ such that $1/\alpha \in \mathbb{N}$, and $\gamma \in (0, 1)$. If τ is an optimal stopping rule from s to π, then there exist γ-valid moves for Crawler such that, whatever (s, α, β)-valid moves Dasher makes,*

$$\mathbb{E}_s[\tau] \leq \frac{1}{\alpha^3} + \frac{2}{\alpha^2\beta\gamma} \sum_{k=1}^{l} \frac{1}{\Phi(D_k)}. \tag{6.1}$$

Proof. We proceed similarly to the proof of Theorem 1.3. Fix a starting state s and $\alpha, \beta \in (0,1]$, $\gamma \in (0,1)$. We describe a strategy for Crawler for which, whatever the strategy of Dasher, any resulting sequence $D_1, \dots, D_l$ satisfies (6.1).

Label the vertices of $V(G)$ as $\{1, \dots, N\}$ such that $y_1 \leq y_2 \leq \ldots \leq y_N$. By Corollary 4.3 we may assume that $s = N$.

Let $v_1 = 1$ and note that v_1 is a halting vertex, i.e. $y_1 = 0$ (since there is always at least one such vertex [13]). Set $C_1 = \{v_1\}$. Dasher can then choose any D_1 that satisfies rules (i) to (iv).

Given C_k, D_k and v_k, if $D_k = V$ then we stop and set $l = k-1$; otherwise we define C_{k+1} and v_{k+1} as follows.

List the vertices of D_k^c as $v_1^k, \dots, v_{n_k}^k$ such that $v_1^k \leq \ldots \leq v_{n_k}^k$. Let C_j^k be the set of vertices of $C_k \cup \{v_1^k, \dots, v_j^k\}$ that are internally connected to v_1.

Choose $m_{k+1} = \min\{m : Q((D_k \cup C_m^k)^c, C_k) \leq \gamma Q(D_k^c, C_k)\}$, or $m_{k+1} = N$ if no such m exists. Let $v_{k+1} = v_{m_{k+1}}^k$ and $C_{k+1} = C_{m_{k+1}}^k$, and $l = \min\{k : m_k = N\} - 1$. Note that $v_{k+1} = \max\{v \in C_{k+1}\}$ (we have $v \geq v_k$ for all $v \in \partial C_k$, otherwise v would have been added to C_k at the previous step; so some vertex larger than v_k is added to C_k in constructing C_{k+1}; and v_{k+1} is the largest vertex added).

By construction, the sequence $(C_1, \dots, C_l)$ satisfies conditions (a) and (b) in the bottleneck sequence game and is therefore a valid strategy. We now show that any sequence $D_1, \dots, D_l$ that can arise when Crawler follows this strategy satisfies (6.1).

Let U_k be the subset of $\{u \in V : u \leq v_k\}$ that is internally connected to v_1. Note that by our construction of C_k, we must have $C_k \subset U_k \subset D_{k-1} \cup C_k$. For $k \leq l$, $s \notin U_k$, so by Lemma 4.2

$$\begin{aligned}
\pi(U_k) &= \sum_{u \in U_k, v \notin U_k} (y_v Q(v,u) - y_u Q(u,v)) \\
&\geq \sum_{u \in U_k, v \notin U_k} y_v Q(v,u) - y_{v_k} Q(U_k, U_k^c) \\
&= \sum_{u \in U_k, v \notin U_k} y_v Q(v,u) - y_{v_k} Q(U_k^c, U_k) \\
&= \sum_{u \in U_k, v \notin U_k} (y_v - y_{v_k}) Q(v,u)
\end{aligned}$$

where we used (1.3) for the penultimate equality. Let $Z_k = (D_k \cup C_{k+1})^c \cup \{v_{k+1}\}$. If $v \notin U_k$ and $\exists u \in U_k$ such that $Q(v,u) > 0$, then $y_v \geq y_{v_k}$; thus, since $C_k \subset U_k$ and $Z_k \subset U_k^c$,

$$\pi(U_k) \geq \sum_{u \in C_k, v \in Z_k} (y_v - y_{v_k}) Q(v,u).$$

Also note that if $v \in Z_k$ and $\exists u \in U_k$ such that $Q(u,v) > 0$, then $y_v \geq y_{v_{k+1}}$. Thus

$$\pi(U_k) \geq \sum_{u \in C_k, v \in Z_k} (y_{v_{k+1}} - y_{v_k}) Q(v,u) = (y_{v_{k+1}} - y_{v_k}) Q(Z_k, C_k). \tag{6.2}$$

By our construction of C_{k+1}, we have

$$Q(Z_k, C_k) = Q((D_k \cup C_{k+1})^c \cup \{v_{k+1}\}, C_k) > \gamma Q(D_k^c, C_k),$$

and by rule (iii) of the bottleneck sequence game applied to (C_k, D_k), with $S = \emptyset$, we have $Q(D_k^c, C_k) \geq \beta Q(D_k^c, D_k)$, so

$$Q(Z_k, C_k) > \beta\gamma Q(D_k^c, D_k).$$

But $U_k \subset D_{k-1} \cup C_k \subset D_k$ (by rule (i) of the bottleneck sequence game), so $\pi(D_k) \geq \pi(U_k)$. Substituting these estimates back into (6.2), we get

$$\pi(D_k) > \beta\gamma(y_{v_{k+1}} - y_{v_k})Q(D_k^c, D_k).$$

Multiplying through by $\pi(D_k^c)$ and rearranging, we obtain the key estimate

$$\pi(D_k^c)(y_{v_{k+1}} - y_{v_k}) < \frac{\pi(D_k^c)\pi(D_k)}{\beta\gamma Q(D_k^c, D_k)} = \frac{1}{\beta\gamma\Phi(D_k)}. \tag{6.3}$$

If $k \leq l$ and $v \in \partial^i D_k$, then by rule (ii), v is α-near to C_k, so by Lemma 6.2 we have

$$y_v \leq \frac{2}{\alpha^2} \max_{u \in C_k} y_u + \frac{1}{\alpha^3}.$$

Suppose $k \leq l$ and take $w \in D_k$. Since, by Corollary 4.3, $\{v : y_v \geq y_w\}$ is a connected set containing both s and w, and $s \in D_k^c$, there must be some vertex $v \in \partial^i D_k$ such that $y_v \geq y_w$. Therefore

$$\max_{w \in D_k} y_w \leq \max_{v \in \partial^i D_k} y_v \leq \frac{2}{\alpha^2} \max_{u \in C_k} y_u + \frac{1}{\alpha^3} = \frac{2}{\alpha^2} y_{v_k} + \frac{1}{\alpha^3}. \tag{6.4}$$

Similarly, using rule (iv) and Lemma 6.2,

$$\max_{w \in D_{l+1}} y_w \leq y_s \leq \frac{2}{\alpha^2} \max_{u \in C_{l+1}} y_u + \frac{1}{\alpha^3} = \frac{2}{\alpha^2} y_{v_{l+1}} + \frac{1}{\alpha^3}. \tag{6.5}$$

Define $q : V \to \mathbb{N}$ by setting $q(v) = k$ if $v \in D_k \setminus D_{k-1}$. Then by (6.4) and (6.5),

$$\mathbb{E}_s[\tau] = \sum_{u \in V} \pi(u) y_u \leq \sum_{u \in V} \pi(u) \max_{w \in D_{q(u)}} y_w \leq \frac{1}{\alpha^3} + \frac{2}{\alpha^2} \sum_{u \in V} \pi(u) y_{v_{q(u)}}.$$

But since $y_{v_1} = y_1 = 0$, we have

$$\begin{aligned}
\sum_{u \in V} \pi(u) y_{v_{q(u)}} &= \sum_{u \in V} \pi(u) \sum_{k=1}^{q(u)-1} (y_{v_{k+1}} - y_{v_k}) \\
&\leq \sum_{u \in V} \pi(u) \sum_{k=1}^{l} (y_{v_{k+1}} - y_{v_k}) \mathbb{1}_{\{u \notin D_k\}} \\
&= \sum_{k=1}^{l} \pi(D_k^c)(y_{v_{k+1}} - y_{v_k})
\end{aligned}$$

where the inequality follows from the fact that if $k \leq q(u) - 1$ then $u \notin D_k$. By (6.3), this is bounded above by $\frac{1}{\beta\gamma} \sum_{k=1}^{l} \frac{1}{\Phi(D_k)}$, and therefore

$$\mathbb{E}_s[\tau] \leq \frac{1}{\alpha^3} + \frac{2}{\alpha^2\beta\gamma} \sum_{k=1}^{l} \frac{1}{\Phi(D_k)}. \qquad \square$$

This completes the proof of Theorem 1.3.

7. Proof of Theorem 1.1

Recall that a Markov chain X on a graph $G = (V, E)$ is ε-uniform if $\varepsilon\pi(x)p_{xy} \leq \pi(u)p_{uv} \leq \pi(x)p_{xy}/\varepsilon$ for all $uv, xy \in E$, and $p_{uv} = 0$ for all $u, v \in V$ with $uv \notin E$.

Our plan for proving Theorem 1.1 is to apply Theorem 1.3, outlining valid moves for Dasher that aim to keep $\sum_k 1/\Phi(D_k)$ small, no matter what valid moves Crawler makes. The nature of rough isometry means that the bounds used in this section will necessarily be somewhat cruder than in previous sections. The following result will be enough to complete the proof.

Proposition 7.1. *Fix a finite connected graph $G = (V(G), E(G))$ and a finite tree $T = (V(T), E(T))$. Suppose that X is an ε-uniform Markov chain on G. Suppose also that $G \simeq_r T$, and G and T both have maximum degree at most Δ. Then for any $s \in V$, there exist $\alpha, \beta, \gamma \in (0, 1)$, depending only on Δ, ε and r, and (s, α, β)-valid moves for Dasher in the bottleneck sequence game, such that whatever γ-valid moves Crawler makes, the resulting sequence $D_1, \ldots, D_l$ satisfies*

$$\sum_{n=1}^{l} |D_n||D_n^c| \lesssim_{\Delta,r} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} |S_j||S_j^c|. \tag{7.1}$$

We delay the proof of Proposition 7.1 for a moment to check that it implies Theorem 1.1.

Proof of Theorem 1.1. Fix any ε-uniform Markov chain Y on T. By (1.2) it suffices to show that $t_{\text{stop}}(X) \asymp_{\Delta,\varepsilon,r} t_{\text{stop}}(Y)$. Note first that since G and T are roughly isometric, we have $|V(G)| \asymp_{\Delta,r} |V(T)|$. It is also easy to check that for any $D \subset V(G)$, we have $\Phi^X(D) \gtrsim_{\Delta,\varepsilon} \frac{|V(G)|}{|D||D^c|}$, and for any $S \subset V(T)$, $\Phi^Y(S) \asymp_{\Delta,\varepsilon} \frac{|V(T)|}{|S||S^c|}$.

We now play the bottleneck sequence game, with Crawler following the strategy from Theorem 1.3 and Dasher following the strategy from Proposition 7.1. Then the resulting sequence $(D_1, \ldots, D_l)$ satisfies

$$t_{\text{stop}}(X) \lesssim_{\Delta,\varepsilon,r} \sum_{k=1}^{l} \frac{1}{\Phi^X(D_k)} \lesssim_{\Delta,\varepsilon} \sum_{k=1}^{l} \frac{|D_k||D_k^c|}{|V(G)|} \lesssim_{\Delta,r} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} \frac{|S_j||S_j^c|}{|V(T)|}.$$

Also, by Proposition 5.1, we have

$$t_{\text{stop}}(X) \gtrsim_{\Delta,\varepsilon,r} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} \frac{|S_j||S_j^c|}{|V(T)|}.$$

But by Corollary 4.5 we have

$$t_{\text{stop}}(Y) \asymp_{\Delta,\varepsilon,r} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} \frac{1}{\Phi^Y(S_j)} \asymp_{\Delta,\varepsilon} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} \frac{|S_j||S_j^c|}{|V(T)|}.$$

Therefore $t_{\text{stop}}(X) \asymp_{\Delta,\varepsilon,r} t_{\text{stop}}(Y)$. □

In order to prove Proposition 7.1, we need to describe a strategy for Dasher. Let $R = 2r^2 - r - 1$.

Given s and Crawler's first move $(C_1, \emptyset)$, let $\tilde{C}_1 = B_G(R, C_1)$, the set of all vertices within distance R of C_1. If $s \in \tilde{C}_1$ then let $D_1 = V(G)$ and $l = 0$ and stop.

Otherwise choose a shortest path $\sigma^1 = (\sigma_0^1, \dots, \sigma_{m_1}^1)$ from C_1 to s. Note that $\sigma_i^1 \in \tilde{C}_n$ for $i \le R$ and σ_i^1 is in the connected component of $\tilde{C}_1^c$ containing s for $i \ge R+1$. Let $D_1 = h_s(\tilde{C}_1) \setminus \{\sigma_1^1, \dots, \sigma_R^1\}$.

We now recursively describe Dasher's nth move for $n \ge 2$; suppose that we are at position (C_n, D_{n-1}), that the previous move was (C_{n-1}, D_{n-1}), and that we have a shortest path $\sigma^{n-1} = (\sigma_0^{n-1}, \dots, \sigma_{m_{n-1}}^{n-1})$ from C_{n-1} to s.

Let $\tilde{C}_n = B_G(R, C_n)$ and $\tilde{C}_{n-1} = B_G(R, C_{n-1})$. If $s \in \tilde{C}_n$ then set $D_n = V(G)$ and $l = n-1$ and stop. Otherwise let $I_n = \max\{i : \sigma_i^{n-1} \in C_n\}$, so that σ_{I_n} is the last vertex in the path σ^n within C_k. Note that (since C_n is connected by rule (a)) if $I_n \le R-1$ then $C_n = C_{n-1} \cup \{\sigma_1^{n-1}, \dots, \sigma_{I_n}^{n-1}\}$. In this case let $\sigma^n = (\sigma_{I_n}^{n-1}, \dots, \sigma_{m_{n-1}}^{n-1})$. On the other hand if $I_n \ge R$ then let σ^n be any shortest path from C_n to s. Finally let $D_n = h_s(\tilde{C}_n) \setminus \{\sigma_1^n, \dots, \sigma_R^n\}$.

We want to show that this strategy is a valid one, and the first thing to check is that D_n has small boundary for every n.

Lemma 7.2. *If $u, v \in \partial D_n$, then $d_G(u,v) \le 2r^2(R+2)+R = 4r^4-2r^3+4r^2+r-1$.*

Proof. Fix a correspondence $\mathcal{C} \subset T \times G$ with stretch at most r. First we show that if $u, v \in \partial \tilde{C}_n$ are in the same component of $\tilde{C}_n^c$, then $d_G(u,v) \le 2r^2(R+2)$. The result will then follow since D_n^c consists of one component of $\tilde{C}_n^c$ together with $\{\sigma_1^n, \dots, \sigma_R^n\}$.

Since u and v are in the same component of $\tilde{C}_n^c$, we may take a path $(u_0, u_1, \dots, u_m)$ from u to v within $\tilde{C}_n^c$. Choose $x, y \in C_n$ such that $d_G(u,x) = R+1$ and $d_G(v,y) = R+1$; since by rule (a) C_n is connected, we may also take a path $x_0, x_1, \dots, x_p$ from x to y within C_n. Clearly $d_G(u_i, x_j) \ge R+1$ for each i and j, and hence $d_T(\mathcal{C}_{T,u_i}, \mathcal{C}_{T,x_j}) \ge (R+2)/r - 1$ for each i and j.

Fix $a \in \mathcal{C}_{T,u}$, $b \in \mathcal{C}_{T,v}$, $t \in \mathcal{C}_{T,x}$ and $t' \in \mathcal{C}_{T,y}$. Note that for any $a_i \in \mathcal{C}_{T,u_i}$ and $t_j \in \mathcal{C}_{T,x_j}$, we have $d_T(a_i, a_{i+1}) \le 2r-1$ and $d_T(t_j, t_{j+1}) \le 2r-1$; therefore we can choose a path γ_1 from a to b and another γ_2 from t to t' such that

$$d_T(\gamma_1, \gamma_2) \ge \frac{R+2}{r} - 1 - 2(r-1) = \frac{2r^2 - r + 1}{r} - 2r + 1 = 1/r > 0.$$

In particular, γ_1 and γ_2 do not intersect.

However, since $d(u,x) = R+1$ and $d(v,y) = R+1$, there exist paths γ_3 from a to t and γ_4 from b to t' of length at most $r(R+2)-1$ each. Note that $\gamma_2 \cup \gamma_3 \cup \gamma_4$ contains a path from a to b, and since T is a tree this path must be γ_1. But γ_1 and γ_2 do not intersect, so γ_1 must be contained in $\gamma_3 \cup \gamma_4$, and therefore be of length at most $2r(R+2)-1$. Therefore $d(a,b) \le 2r(R+2)-1$, so $d(u,v) \le 2r^2(R+2)-1$ as required. □

Lemma 7.3. *There exist $\alpha, \beta \in (0,1)$ (depending only on Δ, ε and r) such that the sequence $D_1, \dots, D_l$ is a valid strategy for Dasher.*

Proof. To check rule (i), clearly $C_n \subset D_n$. Since $C_{n-1} \subset C_n$ by rule (a), to check that $D_{n-1} \subset D_n$ it suffices to check that $\{\sigma_1^n, \dots, \sigma_R^n\} \subset D_{n-1}^c$. If $I_n \le R-1$ this is clear; if $I_n \ge R$, then $d(C_n, s) \le d(\sigma_R^{n-1}, s) \le d(D_{n-1}, s)$, so any shortest path from C_n to s does not intersect D_{n-1}.

For rule (ii), note that $p_{uv} \gtrsim_{\Delta,\varepsilon} 1$ for all $uv \in E(G)$ and $\pi(v) \asymp_{\Delta,\varepsilon} 1/|V(G)|$ for all $v \in V(G)$. It follows that since $\partial^i D_n \subset \tilde{C}_n = B_G(R, C_n)$, there exists α depending only on Δ, ε and r such that $\partial^i D_n$ is α-near to C_n for all n.

For rule (iii), we need to show that there exists $\beta > 0$ such that for any $S \subset D_n^c$ such that $C_n \cup S$ is connected,

$$Q((D_n \cup S)^c, C_n \cup S) \geq \beta Q((D_n \cup S)^c, D_n \cup S).$$

If $S = D_n^c$ then this trivially holds since the right-hand side is 0, so suppose $S \neq D_n^c$. Note that

$$\begin{aligned} Q((D_n \cup S)^c, D_n \cup S) &\leq Q((D_n \cup S)^c, D_n \setminus C_n) + Q((D_n \cup S)^c, C_n \cup S) \\ &\leq Q(D_n^c, D_n) + Q((D_n \cup S)^c, C_n \cup S), \end{aligned}$$

so it suffices to show that for small enough β,

$$Q(D_n^c, D_n) \leq (1/\beta - 1) Q((D_n \cup S)^c, C_n \cup S).$$

But we know from Lemma 7.2 that $Q(D_n^c, D_n) \lesssim_{\Delta,\varepsilon,r} 1/|V(G)|$; and since there is at least one edge between C_n and D_n^c (namely that between σ_0^n and σ_1^n), we have $Q((D_n \cup S)^c, C_n \cup S) \gtrsim_{\Delta,\varepsilon} 1/|V(G)|$. Therefore we may choose β (depending on Δ, ε and r) as required.

Rule (iv) is clearly satisfied. □

Our next lemma shows that our sequence $D_1, \ldots, D_l$ eats up the graph reasonably quickly.

Lemma 7.4. *For any $n, k \geq 1$ such that $n + k \leq l$, we have*

$$d_G(D_n, D_{n+k}^c) \geq k - (4r^4 - 2r^3 + 4r^2 + r - 1).$$

Proof. By our construction of the sequence $D_1, \ldots, D_l$, we have $B_G(k, \{\sigma_0^n\}) \subset D_{n+k}$. Thus $d_G(\sigma_0^n, D_{n+k}^c) \geq k + 1$. By Lemma 7.2, since $\sigma_0^n \sim \sigma_1^n \in \partial D_n$, for any $v \in \partial D_n$ we have

$$d_G(v, D_{n+k}^c) \geq d_G(\sigma_0^n, D_{n+k}^c) - 1 - d_G(\sigma_1^n, v) \geq k + 1 - 1 - (4r^4 - 2r^3 + 4r^2 + r - 1).$$

Since D_n^c is connected, the result follows. □

We now choose the 1-bottleneck sequence on T that we will compare to the sequence $D_1, \ldots, D_l$. Fix a correspondence $\mathcal{C} \subset T \times G$ with stretch at most r.

Take any vertex $v_0 \in C_1$, and then take a simple path $(t_0, \ldots, t_p)$ from $\mathcal{C}_{T,v_0}$ to $\mathcal{C}_{T,s}$. Then for $i = 1, \ldots, p-1$, let $S_i = h_{t_p}(t_i)$, the set of vertices that are separated from t_p by t_i.

It is easy to check that $S_1, \ldots, S_{p-1}$ is a 1-bottleneck sequence for T. If we set $N_k = \min\{i : S_i^c \cap \bigcup_{v \in D_k} \mathcal{C}_{T,v} = \emptyset\}$, then clearly $\bigcup_{v \in D_k} \mathcal{C}_{T,v} \subset S_{N_k}$ and therefore $|D_k| \lesssim_{\Delta,r} |S_{N_k}|$. Let $K = 4r^4 - 2r^3 + 8r^2 - r - 2$.

Lemma 7.5. *For any $n \geq 1$ and $k \geq K$ such that $n + k \leq l$, we have*

$$\bigcup_{v \in D_{n+k}^c} \mathcal{C}_{T,v} \subset S_{N_n}^c$$

and therefore

$$|D_{n+k}^c| \lesssim_{\Delta,r} |S_{N_n}^c|.$$

Proof. Suppose there exists $v \in D^c_{n+k}$ such that $\mathcal{C}_{T,v} \cap S_{N_n} \neq \emptyset$. Since D^c_{n+k} is connected, we can choose a path $v = p_0, \dots, p_m = s$ from v to s within D^c_{k+j}. For each i choose $a_i \in \mathcal{C}_{T,p_i}$, with $a_0 \in S_{N_n}$ and $a_m = t_p$. Then $d_T(a_i, a_{i+1}) \leq 2r-1$, and since T is a tree any path from a_0 to a_m must pass through t_{N_n}. Therefore there exists i such that $d_T(a_i, t_{N_n}) \leq r-1$.

On the other hand, by the definition of N_n, there exists $t \in S_{N_n} \setminus S_{N_n-1}$ and $u \in D_n$ such that $(t,u) \in \mathcal{C}$. Since $D_n \cup \{\sigma^n_1, \dots, \sigma^n_R\}$ is connected, we can choose a path $u = q_0, q_1, \dots, q_{m'} = v_0$ from u to v_0 such that $d_G(q_j, D_n) \leq R$ for all j. For each j take $b_j \in \mathcal{C}_{T,q_j}$ with $b_0 = t \in S_{N_n} \setminus S_{N_n-1}$ and $b'_m = t_0$. Then $d_T(b_j, b_{j+1}) \leq 2r-1$, and since T is a tree any path from b_0 to $b_{m'}$ must pass through t_{N_n}. Therefore there exists j such that $d_T(b_j, t_{N_n}) \leq r-1$.

Putting these two bounds together, we get $d_T(a_i, b_j) \leq 2r-2$, so $d_G(p_i, q_j) \leq r(2r-1)-1$. But $p_i \in D^c_{n+k}$ and $d_G(q_i, D_k) \leq R$, so $d_G(D_k, D^c_{n+k}) \leq r(2r-1)-1+R$. From Lemma 7.4 we see that

$$k - (4r^4 - 2r^3 + 4r^2 + r - 1) \leq d_G(D_k, D^c_{n+k}) \leq r(2r-1) - 1 + R,$$

and rearranging (and recalling that $R = 2r^2 - r - 1$) gives that $k \leq K-1$, from which we deduce the result. □

Next we check that not too many of the sets D_k can give rise to the same value of N_k. This follows easily from Lemma 7.5.

Lemma 7.6. *For any $n \geq 1$ and $k \geq K+1$ such that $n+k \leq l$ we have $N_{n+k} > N_n$.*

Proof. We know from Lemma 7.5 that if $k \geq K$, then

$$\bigcup_{v \in D^c_{n+k}} \mathcal{C}_{T,v} \subset S^c_{N_n}.$$

But if $k \geq K+1$, then D_{n+k} contains at least one element of D^c_{n+K}, and therefore $D_{n+k} \cap S^c_{N_n} \neq \emptyset$. Thus $N_{n+k} > N_n$. □

The next two results are simple technicalities that we will need in our proof of Proposition 7.1.

Lemma 7.7. *For any $j \in \mathbb{N}$,*

$$\sum_{k=1}^{l} |D_k||D^c_k| \leq 2 \sum_{k=1}^{l-j} |D_k||D^c_{k+j}| + \frac{j}{2}|V(G)|^2.$$

Proof. Let $L = \max\{k : |D^c_k| \geq |V(G)|/2\}$. Then

$$\begin{aligned}
\sum_{k=1}^{l-j} |D_k||D^c_{k+j}| &\geq \frac{1}{2} \sum_{k=1}^{L-j} |D_k||D^c_k| + \frac{1}{2} \sum_{k=L+1}^{l-j} |D_{k+j}||D^c_{k+j}| \\
&\geq \frac{1}{2} \sum_{k=1}^{l} |D_k||D^c_k| - \frac{1}{2} \sum_{k=L-j+1}^{L+j} |D_k||D^c_k| \\
&\geq \frac{1}{2} \sum_{k=1}^{l} |D_k||D^c_k| - \frac{1}{2}(2j)\frac{|V(G)|^2}{4}.
\end{aligned}$$

Rearranging gives the result. □

Lemma 7.8. *Provided $|V(T)| \geq 2$, there exists a connected set $A \subset V(T)$ such that A and A^c are both connected with $|A| \geq |V(T)|/(4\Delta)$ and $|A^c| \geq |V(T)|/2$.*

Proof. Take a connected set $B \subset V$ with B^c also connected that maximises $|B||B^c|$. Suppose without loss of generality that $|B| \geq |V(T)|/2$. Choose a vertex $v \in B$ such that v has a neighbour in B^c. Then at least one connected component of $B \setminus \{v\}$ must have size at least $(\frac{|V(T)|}{2} - 1)\frac{1}{\Delta - 1} \geq \frac{|V(T)|}{4\Delta}$. Call this component A. If $|A^c| < |V(T)|/2$, then $|B^c| < |A^c| < |V(T)|/2$ and $|B| > |A| > |V(T)|/2$, so $|A||A^c| > |B||B^c|$, contradicting the maximality of $|B||B^c|$. Therefore $|A^c| \geq |V(T)|/2$. □

Finally we are in a position to complete a proof of Proposition 7.1.

Proof of Proposition 7.1. We have constructed a 1-bottleneck sequence $S_1, \ldots, S_{p-1}$ for T, and a valid strategy for Dasher such that any resulting sequence $D_1, \ldots, D_l$ satisfies $|D_n| \lesssim_{\Delta,r} |S_{N_n}|$ for each n and, by Lemma 7.5, $|D^c_{n+K}| \lesssim_{\Delta,r} |S^c_{N_n}|$ (provided $n + K \leq l$). Applying Lemma 7.7, we obtain

$$\sum_{n=1}^{l} |D_n||D_n^c| \leq 2\sum_{n=1}^{l-K} |D_n||D^c_{n+K}| + \frac{K}{2}|V(G)|^2 \lesssim_{\Delta,r} \sum_{n=1}^{l-K} |S_{N_n}||S^c_{N_n}| + |V(T)|^2.$$

By Lemma 7.6, we have

$$\sum_{n=1}^{l-K} |S_{N_n}||S^c_{N_n}| \leq (K+1)\sum_{n=1}^{p-1} |S_n||S_n^c| \leq (K+1) \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} |S_j||S_j^c|,$$

and by Lemma 7.8 we have

$$\max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} |S_j||S_j^c| \geq \frac{|V(T)|^2}{8\Delta}.$$

Putting these together, we see that

$$\sum_{n=1}^{l} |D_n||D_n^c| \lesssim_{\Delta,r} \max_{(S_1,\ldots,S_m)\in\mathcal{S}_1(T)} \sum_{j=1}^{m} |S_j||S_j^c|$$

as required. □

Acknowledgements. MR was partially supported by EPSRC grant EP/K007440/1. LAB was partially supported by NSERC Discovery Grant 341845.

References

[1] David Aldous. Some inequalities for reversible Markov chains. *J. London Math. Soc*, 25(2):564–576, 1982.

[2] David Aldous and James Allen Fill. Reversible Markov chains and random walks on graphs, 2002. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[3] Itai Benjamini. Coarse geometry and randomness. In *École d'Été de Probabilités de Saint-Flour XLI—2011*, volume 2100 of *Lecture Notes in Math.* Springer, Berlin, 2013.

[4] Jian Ding, James R. Lee, and Yuval Peres. Cover times, blanket times, and majorizing measures. *Annals of Mathematics*, 175(3):1409–1471, 2012.

[5] Jian Ding and Yuval Peres. Sensitivity of mixing times. *Electronic Communications in Probability*, 18:1–6, 2013.

[6] Nikolaos Fountoulakis and Bruce A. Reed. Faster mixing and small bottlenecks. *Probability Theory and Related Fields*, 137(3):475–486, 2007.
[7] Sharad Goel, Ravi Montenegro, and Prasad Tetali. Mixing time bounds via the spectral profile. *Electron. J. Probab*, 11(1):1–26, 2006.
[8] Simon Griffiths, Ross Kang, Roberto Oliveira, and Viresh Patel. Tight inequalities among set hitting times in Markov chains. *Proceedings of the American Mathematical Society*, 142(9):3285–3298, 2014.
[9] Jonathan Hermon. On sensitivity of uniform mixing times. Preprint: `http://arxiv.org/abs/1607.01672`.
[10] Mark Jerrum and Alistair Sinclair. Conductance and the rapid mixing property for markov chains: the approximation of permanent resolved. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 235–244. ACM, 1988.
[11] Gady Kozma. On the precision of the spectral profile. *ALEA*, 3:321–329, 2007.
[12] László Lovász and Ravi Kannan. Faster mixing via average conductance. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 282–287. ACM, 1999.
[13] László Lovász and Peter Winkler. Efficient stopping rules for markov chains. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 76–82. ACM, 1995.
[14] J.W. Moon. Random walks on random trees. *Journal of the Australian Mathematical Society*, 15(01):42–53, 1973.
[15] Ben Morris and Yuval Peres. Evolving sets, mixing and heat kernel bounds. *Probability Theory and Related Fields*, 133(2):245–266, 2005.
[16] Yuval Peres and Perla Sousi. Mixing times are hitting times of large sets. *Journal of Theoretical Probability*, pages 1–32, 2013.

Louigi Addario-Berry, McGill University, Mathematics and Statistics, Burnside Hall, 805 Sherbrooke Street West, Montreal H3A 0B9, Canada

E-mail address: `louigi.addario@mcgill.ca`

Matthew I. Roberts, University of Bath, Department of Mathematical Sciences, Claverton Down, Bath BA1 1AR, UK

E-mail address: `mattiroberts@gmail.com`